

**IN THE UNITED STATE PATENT AND TRADEMARK OFFICE
APPLICATION FOR UNITED STATES LETTERS PATENT**

**DISTRIBUTED EXTRACT, TRANSFER, AND LOAD (ETL)
COMPUTER METHOD**

By:

**Duane L. Porter
2813 NW Westbrook Circle
Blue Springs, Missouri 64015
Citizenship: United States**

**David E. Swanholm
16415 Nall Avenue
Stilwell, Kansas 66085
Citizenship: United States**

5 TITLE

6 Distributed Extract, Transfer, and Load (ETL) Computer Method

CROSS-REFERENCE TO RELATED APPLICATIONS

8 Not applicable.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

11 Not applicable.

REFERENCE TO A MICROFICHE APPENDIX

13 Not applicable.

FIELD OF THE INVENTION

15 [0001] The invention is a distributed extract, transfer, and load (ETL) computer method
16 (hereinafter referred to as Distributed ETL) for linking together multiple information domains
17 residing across an enterprise-wide computing environment in order to share corporate
18 information, while reducing the time required to deliver that information. The invention
19 overcomes the many-to-many relationships that exist between information sources and targets
20 residing within each domain by implementing a central router to deliver information to and from
21 these domains.

BACKGROUND OF THE INVENTION

23 [0002] Today, better integration of business information systems throughout the enterprise
24 and incorporating more recent database management technologies results in real managerial
25 competitive advantages. Information to support management decision making typically is
26 extracted from operations systems and loaded into a data warehouse. Information as used herein
27 includes data internal to a corporation that is generated in the day-to-day business operations of
28 the corporation, aggregations and summaries of the internal data, and external information such

5 as demographic and geospatial data that can be related to the internal data. A data warehouse is
6 a large storage facility within the information technology, IT, system that typically holds the
7 current and historical information across an entire enterprise. It most likely consists of numerous
8 databases and may exist on multiple platforms (operating systems). A database can be defined as
9 a collection of information organized in such a way that a computer program can quickly select
10 desired pieces of data. One can think of a database as an electronic filing system.

11 [0003] In the past, it was only viewed as necessary to capture periodic data extracts supplied
12 by operational computing systems to update the data warehouse on a weekly or monthly basis.
13 Strategic dependence on data warehousing was viewed as long-term and historical, not as a real
14 time activity. In contrast, today's business environment is worldwide and does not stop at five
15 o'clock P.M. For example, the Internet allows commerce to continue 24 hours a day, seven days
16 a week. As a result, customer service support and normal telephony operations within many
17 business organizations likewise operate around the clock.

18 [0004] Furthermore, marketing practices have evolved from mass media advertising to more
19 personalized target marketing, which cuts expenses while providing customized treatment. To
20 enable customized marketing, companies must know and understand their customers. Much of
21 this understanding is made possible through groups such as Decision Support Services (DSS),
22 the business group within an enterprise that has traditionally supported data warehousing. To be
23 effective in supporting the extended operating and marketing efforts of an enterprise, information
24 should be processed and pushed back from DSS into the customer contact areas such as customer
25 service support, telemarketing and advertising campaign management as rapidly as possible.
26 Thus, it is preferred that DSS be capable of receiving, processing, and providing information on
27 an ongoing, near real time basis.

5 [0005] This invention, referred to as Distributed ETL, improves on current DSS services by
6 providing the advantages of scalability of an extract, transfer and load (ETL) application and the
7 near real time event processing of an enterprise application integration (EAI) application without
8 the disadvantages associated with each.

9 **SUMMARY OF THE INVENTION**

10 [0006] The present invention is a distributed extract, transform and load (ETL) method for
11 delivering information within a computing environment, comprising extracting information from
12 an information source and transforming the extracted information. The transformed information
13 is isolated by wrapping the transformed information into a message envelope having a standard
14 format. The message envelope is routed to at least one information target, unwrapped to reveal
15 the received information, preferably transformed again, and loaded into the information target.
16 The extraction, transformation, and adaptation steps on the source side are isolated from the
17 routing step such that the extraction, transformation, and adaptation steps on the source side may
18 be executed simultaneously for a plurality of information sources distributed across the
19 computing environment to produce a plurality of message envelopes. The routing, unwrapping,
20 mapping, transformation, and loading steps on the target side are repeated for each of the
21 plurality of message envelopes.

22 **DESCRIPTION OF DRAWINGS**

23 [0007] Figure 1 is a diagrammatic representation of business information domains.
24 [0008] Figure 2 is a diagrammatic representation of a basic hub and spoke architecture.
25 [0009] Figure 3 depicts the Distributed ETL architecture of the current invention.
26 [0010] Figure 4 is a diagrammatic representation of a staging area.
27 [0011] Figure 5 is a detailed diagrammatic representation of the router components.

5 [0012] Figure 6 depicts the utilization of the meta data repository by the adapter component
6 of the Distributed ETL application and a parsing of the message envelope that is distributed by
7 the router component.

8 [0013] Figure 7 is a more detailed diagram of the source and target spoke components.

9 [0014] Figure 8 is a diagrammatic representation of information flow with the Distributed
10 ETL implemented on a single platform.

11 [0015] Figure 9 is a diagrammatic representation of a preferred embodiment of the
12 Distributed ETL routing across multiple platforms.

13 [0016] Figure 10 depicts a general overview of how the Distributed ETL architecture
14 integrates the various business enterprise information domains.

15 **DETAILED DESCRIPTION OF THE INVENTION**

16 [0017] Within the enterprise, information takes many forms depending on its usage. In other
17 words, the business usage of the information drives its resultant structure. This information can
18 be grouped into information domains according to its business usage or function as shown
19 diagrammatically in Figure 1. The information domains include operations 10 used for daily
20 operation information such as on-line transaction processing (OLTP); exploration zones 120
21 which includes data mining; data warehouse 40 which is the consolidated information store; and
22 summarized views 150 used for predefined queries and on-line analytical processing (OLAP)
23 reports.

24 [0018] As shown in Figure 1, the information domains overlap, indicating that the
25 boundaries between business information domains are not absolutely crisp. Some warehouse
26 information can be summarized, a characteristic of OLAP summary reporting, while OLAP tools
27 sometimes drill down into detailed warehouse tables. Because of the differing data structures

5 required to support the differing business uses for the data, the same information may be used for
6 multiple business functions. Therefore, the information domains are interdependent upon each
7 other. The relationship between information domains can become a complex web linking
8 information sources and information targets. These relationships can be intermixed and
9 multidirectional, thus defining a many-to-many relationship between business information
10 domains.

11 [0019] Current enterprise application integration (EAI) applications, such as IBM's MQ
12 Series, Tibco's Rendezvous or Saga's Saga Vista, are available to coordinate information flow
13 between information domains to facilitate business processes. These traditional EAI systems are
14 synchronous, transaction driven systems that operate in real time or near-real time, but have
15 certain other disadvantages (e.g., scalability limitations). An alternative tool is a standard
16 extract, transform and load (ETL) application, which is an asynchronous, batch driven system
17 that can be scaled up but typically suffers from delayed event processing.

18 [0020] Traditional EAI applications direct the flow of information between domains through
19 a center point called a hub, or more commonly a message broker to manage many-to-many
20 relationships in a scalable manner. The hub and spoke architecture is applicable to managing
21 ETL needs that exist between information domains as well. Figure 2 is a diagrammatic
22 representation of a basic hub and spoke architecture. The staging hub 70 is at the center of the
23 hub and spoke figure with connecting spokes 2, 3, 4, and 5 to the various information domains,
24 namely operations 10, exploration zones 120, data warehouse 40, and summarized views 150,
25 respectively.

26 [0021] In contrast to the traditional EAI and standard ETL applications which rely on the
27 basic hub and spoke architecture described above, this invention is a distributed extract, transfer,

5 and load (ETL) computer application to address the need for linking multiple information
6 sources and targets together to deliver shared corporate information, while reducing the time
7 required to deliver that information. The Distributed ETL overcomes the many-to-many
8 relationships that exist between information sources and targets by implementing a spartan
9 central router passing information to and from these domains through “thick spokes.”

10 [0022] Figure 3 depicts the Distributed ETL architecture of the current invention.
11 Distributed ETL is comprised of two parts. First, a set of distributed “thick spokes” 20, 50, 100
12 and 130 which provide pre-routing information processing specific to sources connected by input
13 spokes and post-routing information processing specific to targets connected by output spokes.
14 Second, a spartan router component 75 serves as a central core by connecting the input and
15 output thick spokes. Preferably, all shared information in the enterprise will flow through the
16 router component.

17 [0023] The “thick” spokes 20, 50, 100 and 130 are used to manage specific source/target
18 information processing (also referred to as staging processes, which comprise extraction,
19 transformation, adaptation, and loading functions) as may be required prior to and/or after
20 routing of the information. In contrast, the router component 75 is spartan in that it primarily
21 manages delivery (i.e., routing) the information and does not perform processing steps thereon.
22 Staging processes that are specific to a source or target are pushed out closer to where the spoke
23 connects to that source or target, in contrast to being performed by the router component. By
24 isolating the router component to purely delivery functions, the entire architecture is streamlined,
25 enhancing its speed, efficiency, parallelism, and scalability. In sum, the architecture of the
26 present invention comprises: 1) the staging processes to manipulate the information (in the

5 “thick” spokes), and 2) the spartan router component to deliver the information. Therefore, the
6 overall architecture can best be referred to as a Distributed ETL (extract, transform and load).

7 [0024] The staging processes, or the “thick” spokes of the Distributed ETL, are comprised of
8 several software components to process information as necessary prior to and/or after routing it.
9 Referring to Figure 4, a Distributed ETL architecture is shown wherein information is being
10 routed from any one of two information sources 12 and 42 to any one of two information targets
11 122 and 152. Any operable number of information sources and targets can be linked using the
12 Distributed ETL of this invention, and the number of sources and targets shown in Fig. 4 was
13 chosen for convenience purposes only. On the information source side of the router component
14 75 (i.e., the source “thick” spoke), the Distributed ETL comprises extract components 22 and 52;
15 transform components 26 and 56; and source adapter components 28 and 58. On the information
16 target side of the router component 75 (i.e., the target “thick” spoke), the Distributed ETL
17 comprises target adapter components 102 and 132; transform components 106 and 136; and load
18 components 116 and 146. Additional adapter components 24 and 107 may optionally be used on
19 the source spoke and the target spoke, respectively. Combined, these components insolate the
20 information sources and targets from the router component and provide a standard interface for
21 the router.

22 [0025] The first step, beginning at an information source locations 12 or 42 involves the
23 extraction of information by the extract component 22 or 52. The extract component is a
24 software program that functions as a connector to an external information source for retrieving
25 information from the source. Information can be pulled or pushed from the source, and
26 preferably the source pushes the information to the Distributed ETL environment. In most cases,
27 it is desirable to retrieve only updates, revisions, changes and the like to the information

5 ("information changes"), as this limits the amount of information that needs to be delivered to
6 the targets. Methods to extract the information changes from the source include:

- 7 1. A full extract of the source information is created, which is then compared against a
8 previous full extract that was created at an earlier point in time. A detailed
9 comparison between the two extracts generates the list of information changes.
- 10 2. A relational database management system (RDBMS) creates a log of all changes
11 made to the database. This log can be read periodically to obtain the changes made to
12 the source information.
- 13 3. The RDBMS can utilize what are called "triggers" that perform actions when certain
14 data tables are affected and identified conditions occur, pushing the information
15 changes to a table or queue that may be accessed for processing. The principal
16 disadvantage to this approach is that it often loses transactional integrity. For
17 example, where a single transaction for a customer spans several related tables,
18 triggers looking for changes to only a single table may lose sight of the relationships
19 needed to define all of the information changes.
- 20 4. The RDBMS can use a trigger to execute a "stored procedure" to capture information
21 changes. In this case, the trigger executes when it senses a key relation change, and
22 initiates a programmed stored procedure that gathers all information relevant to a
23 logical transaction, based on the business rules for relationships contained in the
24 referential integrity of the database design. The stored procedure then passes the
25 information changes to the Distributed ETL with full transactional integrity.

26 [0026] The extracted information is then passed from extract component 22 or 52 to the
27 transform component 26 or 56, respectively. Operational systems are the primary information

5 sources in an enterprise and typically employ highly normalized data structures that optimize
6 OLTP performance critical to real time, customer facing processes. Normalized data structures
7 tend to use multiple tables to minimize recurring data forms and data redundancy. Data
8 redundancy and recurring formats within records/rows can cause severe performance penalties in
9 an environment having a high number of information updates. The fragmentation of normalized
10 design, useful in operational systems, creates problems when shipping the information in a
11 normalized format to remote targets. The target must know the business rules, or referential
12 integrity, that logically ties the several table structures together as an integrated, functioning
13 whole. In order to pass the normalized data in a message format, each piece of information must
14 be tagged in a way that relates it to the other associated information so they can all be delivered
15 together and assembled according to the master blueprint to a state of referential integrity, which
16 is very impractical. Thus, in order to implement message based transport, denormalization of the
17 data to some extent is required so that sufficient information is contained in each message to
18 allow interested targets to easily identify and subscribe to events they wish to receive. This is the
19 essence of packet based messaging, i.e., that each packet is self-addressed by the information
20 contained within it. The denormalization process required to create these packets relies on
21 transformation steps, with full knowledge of the source referential integrity rules (i.e., business
22 rules governing the information structure).

23 [0027] The transform component's function is to perform transformations on the extracted
24 information. As described above, transformation refers to the application of specific business
25 rules governing the structure of the information. For example, captured information such as the
26 price of a certain product and the quantity sold to a customer may need to be transformed. The
27 desired form of this information might be the extended price. By applying the business rule for

5 extended price, e.g., product price multiplied by quantity sold, the information has been
6 transformed to the desired structure. Another example of information transformation is where a
7 source captures customer names in the form of full name including the middle initial, e.g., John
8 D. Smith. But, the desired form of this information for one specific target could be first and
9 middle initials followed by the last name, e.g., J. D. Smith. Additionally, a different target may
10 desire the same information in yet another format, such as last name first, followed by first name,
11 followed by middle initial, e.g., Smith, John D. The last two examples demonstrate how two
12 different targets can access the same source information and transform it in different ways that
13 are acceptable to their business needs.

14 [0028] Transformations are used whenever the format of the information at the target, as
15 defined by the business needs, does not match the format of the information extracted or arriving
16 from the source. Stated alternatively, transformations are used to apply business rules that are
17 common to the enterprise on the source side, and specific to the target on the target side.
18 Typically, a source-side transformation is performed in order to facilitate packet based
19 messaging, as discussed previously. Also, a target-side transformation typically is performed in
20 order to facilitate the incorporation of information from multiple sources into a single target. In
21 other words, a benefit of information targets is that they consolidate information from multiple
22 sources, which frequently involves transformation. Even where a target receives information
23 only from a single source, typically the target adds value by creating different perspectives of
24 that information through aggregation and summarization, which also involve transformation.
25 Where the target only displays detailed data without further transformation, such information
26 may be read (i.e., queried) directly from an enterprise data warehouse, which typically has all
27 source rules already applied. However, a data warehouse typically uses a normalized data

5 structure, and thus a transformation may be required in SQL queries to return information in a
6 meaningful format.

7 [0029] The transformed information is now ready to be delivered to the target but must first
8 pass through the source adapter component 28 or 58. The source adapter component functions
9 asynchronously as a connector to the router component 75 thereby isolating transformations from
10 the router component and ensuring that the router component's primary function is delivery of
11 the information (in contrast to the pre-routing information processing in the "thick" source
12 spokes and the post-routing information processing in the "thick" target spokes). Herein lies
13 (both for the source adapter and as discussed below, for the target adapter) the adapter
14 component's primary function, to standardize and simplify the connection to the router
15 component. The source adapter component 28 or 58 is implemented in software and wraps the
16 transformed information in a standard message format, preferably a structured event format such
17 as CosNotification, which is part of the publicly available CORBA standard. The
18 CosNotification structured event is supported by various application server platforms such as
19 WebLogic available from BEA Systems, Inc. Hereinafter, the adapted information is referred to
20 as a message envelope.

21 [0030] The message envelope is sent to a message queue by the adapter to await distribution
22 via the router component 75. The source adapter must be intimately aware of content of the
23 extracted and transformed information given that a single source may provide multiple types of
24 information such as an addition, a revision, or a deletion of information. Because some targets
25 may wish to subscribe to, or receive, only some of these information types, the source adapter
26 must provide the specific information type for each message envelope sent to the router. To do
27 this, the source adapter must have access to the business rules that define how to identify each

5 type of information. For example, a certain code or date value may differentiate between an
6 addition and a revision. Once the information type is identified, the source adapter tags the
7 message envelope header with the proper message name (also referred to as the source ID) that
8 will then be used for routing to the appropriate subscribing targets.

9 [0031] Meta data is used for control and management of information within the Distributed
10 ETL application. Meta data is defined as “data about data” or facts about the information
11 architecture within an organization. Meta data describes how, when, and by whom a particular
12 set of information was collected, and how the information is formatted. Meta data is essential for
13 understanding information stored in data warehouses, and the meta data is stored in a meta data
14 repository. An example of meta data stored in the repository is transformation rules that describe
15 how to move information from a transactional source database to a multidimensional target
16 database and the aggregations and calculations that are applied to the information coincident to
17 such a move. Another example of meta data in the repository is business rules (such as the
18 definition of revenue or a common list of product codes) that multiple software components may
19 want to access. Another example is when moving information from a source to a target, the meta
20 data is checked to determine the appropriate characteristics of the data. For example, where a
21 numeric field is required, an alpha numeric information piece would be rejected. When the
22 adapter component detects such an error, the adapter may either route the information to an error
23 queue rather than sending it on to the router component, or alternatively correct the information
24 on the fly according to predetermined business rules. In sum, meta data provides transparent
25 access to information.

26 [0032] Because the information processing within the thick spokes of the Distributed ETL
27 runs in parallel, synchronization is important and can be accomplished by tracking the

5 information with meta data. The source adapter component utilizes a meta data repository to
6 provide a “road map” and rules for the information that it processes. The meta data stored within
7 the meta data repository can be referred to as the content definition. Analogous to a COBOL
8 copybook in a COBOL program, the content definition specifies the structure and naming
9 conventions for the content, that is how specific fields are names, field sizes, formats, etc. within
10 the record content. The key associated with this content definition that defines where in the meta
11 data repository the full content definition can be found is referred to as the content definition ID.
12 These are the keys to the meta data and only the keys need to be passed through the router
13 component rather than the whole content definition.

14 [0033] In a preferred embodiment, the target knows the source that it will be subscribing to
15 in advance and simply references the content definition from the metadata repository when the
16 target adapter is constructed. Thus, no content definition IDs (i.e., keys) need to be passed at all.
17 Alternatively, where the content of each message may rely on a different document type
18 definition (DTD) format, for example source information comprising extended markup language
19 (XML) files, the key for the DTD may be sent with the message for retrieval on the target side
20 rather than shipping the entire DTD with each message envelope. This significantly reduces the
21 data transmitted and increases efficiency.

22 [0034] The utilization of the meta data repository by the adapter components of the
23 Distributed ETL application and a parsing of the message envelope that is distributed by the
24 router component are depicted in Figures 6. The source adapter component 28 pulls the content
25 definition 61 from the meta data repository 60. The adapter then uses the content definition 61 to
26 help process the transformed information (referred to as content 29) passed to it from the
27 transform component as previously described in Figure 4. The content is the information

5 extracted from the source, expanded and/or modified by the transform, and its definition is stored
6 in the metadata store. Preferably, the content definition is defined and stored prior to the
7 processing of any information by the adapter. The adapter creates the message envelope (using a
8 structured event format) by performing simple mapping of the content to a uniform format and
9 by building header information such as the message name and source ID, which is used to route
10 the message envelope.

11 [0035] The adapter component attaches a header to the content creating a message envelope
12 30. The header includes a source ID 31 used to route the content to the appropriate targets by the
13 router component. The message envelope 30 may also contain the content definition ID 32 and
14 preferably will include the content 33. The router component delivers the message envelope 30
15 by utilizing a router table 86 as shown in Figure 5 and discussed below. The router table 86 is
16 stored in local memory to do an address lookup to find the subscribing target 35 and match its
17 source ID 34 with the source ID 31 in the message envelope 30.

18 [0036] Referring to Figure 4 on the exiting side of the routing function (i.e., the thick target
19 spoke), near the target or targets, additional adapters 102 and/or 132 implemented in software
20 read the router component output queue and remove the envelope from the distributed
21 information. Referring to Figure 6 the adapter 102 retrieves the content definition 62 from the
22 meta data repository 60 by reference to the content definition ID 32 of the retrieved message
23 envelope 30. Referring to Figure 4 transformations can be made in the target thick spoke (i.e.,
24 post-delivery by the router component), if necessary, according to the target requirements, with
25 an additional transform component 106 or 136. The information is then passed to the respective
26 load component 116 or 146. The load component preferably uses an Application Programming
27 Interface (API) software function to load the information into the target 122 and 152. An API is

5 useful to abstract the load operation from the specific characteristics of a given physical
6 datastore/file. In a typical embodiment, the target is a data warehouse and/or a data mart.

7 [0037] Multiple adapters are possible in the source and target thick spokes, as shown in
8 Figure 4. On the source side, additional source adapter 24 is shown between the extract
9 component 52 and transform component 56. On the target side, additional target adapter 107 is
10 shown between transform component 136 and load component 146. These additional adapters
11 allow for different technologies to coexist without requiring extensive and costly integration.

12 For example, an additional adapter may be used on the source side for transporting an extracted
13 file to a different platform for transformation. Another example, an additional adapter may be
14 used to enable a transactional system to flow events directly into a transformation process in
15 near-real-time (shown as process 7 in Figure 7).

16 [0038] The router component 75 is an important component in the Distributed ETL, because
17 all information sources and targets are connected through it. The key function of the router
18 component is to successfully route each message envelope from its publishing source to a single
19 subscribing target or multiple subscribing targets. It is comprised of multiple software
20 components that are message based, reading input from and writing output to message queues,
21 and should be generalized to the extreme. Ideally, routing should be easy to change and require
22 a minimal time to execute.

23 [0039] The router component 75 uses memory-based messaging to transfer information
24 between publishing sources and subscribing targets. More specifically, messaging occurs on the
25 source side between the source adapter and the router and on the target side between the router
26 and the target adapter. In summary, information is extracted from a source and transformed.
27 The adapter component creates the message envelope and publishes the source message to the

5 router via a message queuing mechanism. The router subscribes to all source queues, thereby
6 receiving the message envelopes to be routed. The router performs a lookup in the routing tables
7 using the source ID as a key. The router publishes the message envelope to one or more target
8 destination queues. The target adapter subscribes to a message queuing mechanism, unwraps the
9 message envelope, and passes the information to a target transformation component. The
10 information is then passed to a load component where it is written directly to the ultimate target.
11 Linking distributed sources or targets with messaging functionality enables a minimum handling
12 time per unit of information. Pairing transactional information extraction and loading with
13 messaging enables near real time processing of extraction and loading functions.

14 [0040] Guaranteed message delivery as a part of the publish and subscribe functionality is
15 important to ensure system reliability and verify that during exchange of messages no
16 information packets are lost. When the network is available, messages will be delivered
17 immediately. When the network is down, messages are still guaranteed to be delivered,
18 eventually, but with some latency. In the latter case, middleware guarantees that messages are
19 delivered as long as the network becomes available within a specified time period. In addition,
20 messaging should provide the assurance of non-duplicate delivery, meaning the messages are
21 only delivered once.

22 [0041] Messaging utilizes a queue system, which allows for a plurality of information
23 sources to simultaneously publish to a router component inbox queue, preferably where each of
24 the sources is capable of publishing independently of the others. A source (via the extract,
25 transform, and adapt components discussed previously) can publish to an inbox queue while the
26 router component is subscribing to that inbox queue. Unlike a sequential file, the router
27 component can read message envelopes from the front of the inbox queue at the same time that

5 the source is publishing new message envelopes into the back of the inbox queue. This allows
6 faster aggregate throughput than a traditional batch environment. A router component can then
7 publish a message envelope to a subscribing target's queue, which the target can access (as
8 discussed previously with regard to the target adapter) and store for later retrieval matching its
9 business cycle needs.

10 [0042] Performing its functionality in memory keeps the router component robust, as
11 memory is hundreds to thousands of times faster to access than disk. Thus, it is preferred to use
12 routing tables that can be maintained in memory for the control of routing destinations. The
13 routing tables are backed up by persistent storage such that the tables may be reloaded to
14 memory when the power is restored after a failure, or to update the content of the table when
15 target subscriptions are changed.

16 [0043] Figure 5 is a diagrammatic representation of the router components and functionality
17 comprising the Distributed ETL architecture. The router component 75 receives a message
18 envelope 82 from any one of multiple sources 12, 13, 14, 182, 183 or 184 through the thick
19 source spokes and the inbox queue discussed previously. The message envelope is then passed
20 to an address lookup 84 where a publish/subscribe cross-reference table 86 (also referred to as a
21 router table) is accessed to determine where the message envelope is to be sent. The header on
22 the message envelope identifying the publisher (i.e., source) corresponds to a list of subscribers
23 (i.e., targets) in the cross-reference table 86 so the router component can determine the routing
24 destination. The message envelope is then sent from Address Lookup 84 to Envelope Re-
25 Address 94 where the target address is added to the message envelope. In an alternative
26 embodiment, the source adapter automatically sends the message envelope to the router, which
27 in turn routes the message envelope based upon the originating source (via the lookup table) and

5 with the original source name intact. In other words, the target address is found in the lookup
6 table and is not placed upon the message itself. The message envelope is then sent out to the
7 appropriate target via Envelope Transmit 96. The output from the router component is written to
8 an output queue that will be read by an adapter component on the target side of the distribution
9 function. Persistent Publish/Subscribe Tables 90 are updated through a Publish/Subscribe
10 Manager 92, which provides a Graphic User Interface (GUI) configured for table management.
11 These updates are sent to the Memory Refresh 88 to provide current entries in Publish/Subscribe
12 cross-reference table 86. Addresses contained within the cross-reference table are stored in local
13 memory within the router to quicken address lookup time. The Error Handle and Administration
14 97 handles errors by sending an error message, Alert Message Handling 98 and storing the errors
15 in a Routing Error Cache/Queue 99.

16 [0044] Downstream of the routing function another adapter 102 reads the message envelope
17 30. It uses the content definition 62 retrieved from the meta data repository 60 to extract the
18 content 101 from the distributed information, as described previously. The information is
19 subsequently processed in the thick target spoke on its way to the appropriate target as described
20 previously in relation to Figure 4.

21 [0045] Figure 7 is a more detailed diagram of the components comprising the thick source
22 and target spokes. As described previously, information is extracted from the information source
23 12 or process 7 by the extract component 22. The extracted information is passed to the
24 transform component 26 where the information is transformed. The extracted, transformed
25 information is sent through the adapter component 28 where it is wrapped in a message envelope
26 and sent to the inbox queue for delivery by the router component 75.

5 [0046] On the output side of the router component 75, the adapter component 102 reads the
6 information from a message queue and takes the information out of the message envelope
7 (removes the header). In a preferred embodiment, the source information (including the type of
8 information such as an addition, deletion, etc.) contained within the message envelope is specific
9 enough such that the router component can route the message envelope only to targets desiring
10 (i.e., subscribing) such information, thereby eliminating the need for filtering. In a non-preferred
11 embodiment, the adapter component 102 passes the information through an optional filter
12 component 104. The filter component's function is to only accept information required by the
13 target and to reject non-conforming information. The filtered information may now undergo
14 more transformations by the transform component 106 that can be applied at the event level, i.e.,
15 one message at a time. In other words, all information need not be made available at the same
16 time such as sorting a group of records, all of which must be made available in order to perform
17 the search. In contrast for example, a trunk code could be used in a message to perform a lookup
18 and the trunk location returned to update the same record. Where the information has been
19 acquired continuously over a period of time, the information in a preferred embodiment is passed
20 to the time-period aggregator component 108. Here the information can be accumulated in an
21 "active" partition or file and later retrieved by toggling the "active" partition 110 to an "inactive"
22 state 112, allowing it to be read into the target transformation. The information can be held to
23 accommodate a business time cycle, e.g., weekly or monthly reporting. When the partition is
24 toggled to an inactivate state 112, the accumulated information can now be passed on for
25 additional transformations by an additional target transform component 114 where necessary.
26 This level of aggregation (in contrast to a time period aggregation) accommodates business
27 aggregations, e.g. by region, product, salesperson, etc. as well as other transformation logic such

5 as any matching operations if multiple information sources are received for consolidation. The
6 information is then passed to the loading component 116 where the information is finally loaded
7 into the target 122 typically as a batch process.

8 [0047] Figure 8 is a diagrammatic representation of information flow on a single platform of
9 a Distributed ETL. A single platform means that all the information sources and targets are
10 routed through one router component. The single platform could represent a small company's
11 information system or a portion of a larger company's information system. There can be
12 multiple sources 12, 13 and 14 (connected to the router component by thick source spoke 330)
13 and multiple targets 122, 123 and 124 (connected to the router component by thick target spoke
14 335). Note that targets can also be sources 182, 183 and 184 known as recursive targets, both
15 receiving information through a thick target spoke 340 and sending information through a thick
16 source spoke 340. A recursive target receives information from the router, adds some unique
17 transformation value to the information and sends it back through the router to another specific
18 target or targets. A recursive spoke does not send out the same information it receives.
19 Information flow begins at the information source 12, 13 or 14 travels through thick source
20 spoke 330 to the router component 75 and then on through thick target spoke 335 or 340 to either
21 the information target 122, 123 or 124 or to the recursive target 182, 183 or 184, respectively.
22 The uniquely transformed values from the recursive targets are then routed back through thick
23 source spoke 340 to the router component 75 and on through thick target spoke 335 to the
24 information target 122, 123 or 124.

25 [0048] In contrast to the single platform structure of Figure 8, Figure 9 is a diagrammatic
26 representation of a preferred embodiment of the Distributed ETL routing environment across
27 multiple platforms. A plurality of sources and targets 12-14, 42-47, 182-184, 122-124, 152-157,

5 282-287 within an enterprise can be logically grouped into multiple platforms as shown with
6 platform A 11, platform B 41, and platform C 121. By introducing multiple routers 75, 76 and
7 77 into the overall architecture which deliver message envelopes to one another a represented by
8 arrows 300, 305, and 310, better utilization of bandwidth between devices can further enhance
9 efficiency as well as make the design easier to scale up to very large business intelligence
10 integration systems.

11 [0049] Integration of the various business enterprise information domains from Figure 3 with
12 the Distributed ETL architecture of Figure 9 is depicted in Figure 10. The information domains
13 10, 40, 120 and 150 link to a logical hub 70 to resolve their many-to-many relationships.
14 Platforms A 11, B 41, and C 121 show the physical means by which this is accomplished within
15 the logical hub. The physical platforms contain multiple source and target spokes, which could
16 connect to any of the information domains. Information is shared between physical platforms
17 through the routers that exist on each platform 75, 76, 77. Figure 10 shows how the physical
18 solution fits into the larger logical/conceptual picture.

19 [0050] In a typical embodiment, the information target comprises a data warehouse and a
20 data mart. Since it can be quite cumbersome to access information from the data warehouse
21 because of the large amounts of data and typically normalized data structures, a common IT
22 design feature is to provide smaller information storage places that focus on a particular subject
23 or department, called data marts. Historically, these can either be dependent data marts, which
24 are subsets of the data warehouse, or independent data marts, which load information directly
25 from operations (point to point). Data marts function to provide easier access to stored
26 information. The data marts are highly indexed and optimized for query. Where one knows in
27 advance the questions to be asked, data marts can be structured accordingly.

5 [0051] Legacy systems utilizing dependent data marts tend to be slow because it takes time
6 to first load the data warehouse and then initiate another process to load the data marts. To
7 reduce latency, an efficiency feature of the Distributed ETL invention loads information directly
8 into ‘co-dependent’ data marts instead of the data warehouse. Co-dependent data marts are built
9 using information that is dependent upon the same rules as the data warehouse but does not
10 require the latency involved in writing to a data warehouse structure followed by extraction,
11 transport and/or transformation and a final write to the dependent data mart.

CONFIDENTIAL - DRAFT